

1           What is claimed is:

2

3       1.     In a file server system having a clock for producing a clock time and a processor  
4       for servicing client requests for access to a file, the processor having a timer for  
5       measuring a time interval, a method comprising:

6           the processor obtaining the clock time from the clock, and beginning  
7       measurement of the time interval with the timer, and

8           the processor responding to a request from a client for an asynchronous write to  
9       the file by performing an asynchronous write operation with respect to the file, and  
10       determining a file-modification time that is a function of the clock time having been  
11       obtained from the clock and the time interval measured by the timer, the file-modification  
12       time indicating a time of modification of the file by the asynchronous write operation.

13

14       2.     The method as claimed in claim 1, wherein the file-modification time is a sum of  
15       the clock time having been obtained from the clock and the time interval measured by the  
16       timer.

17

18       3.     The method as claimed in claim 1, which includes the processor acknowledging  
19       the request from the client for an asynchronous write to the file by returning to the client  
20       the file-modification time.

21

22       4.     The method as claimed in claim 1, which further includes the processor receiving  
23       an updated value for the file-modification time after the processor has determined a value

1 for the file-modification time, the processor comparing the updated value to the value that  
2 the processor has determined for the file-modification time, and upon finding that the  
3 updated value for the file-modification time is greater than the value that the processor  
4 has determined for the file-modification time, then the processor resetting the timer and  
5 using the updated value for the file-modification time in lieu of the clock time obtained  
6 from the clock.

7  
8 5. The method as claimed in claim 4, wherein the processor stores the clock time  
9 having been obtained from the clock in a memory location local to the processor, and the  
10 processor uses the updated value for the file-modification time in lieu of the clock time  
11 obtained from the clock by replacing the clock time having been obtained from the clock  
12 and stored in the memory location local to the processor with the updated value for the  
13 file-modification time.

14  
15 6. The method as claimed in claim 1, which further includes the processor receiving  
16 an updated value for the file-modification time after the processor has determined a value  
17 for the file-modification time, the processor comparing the updated value for the file-  
18 modification time to the value that the processor has determined for the file-modification  
19 time, and upon finding the updated value for the file-modification time is less than the  
20 value that the processor has determined for the file-modification time, then the processor  
21 ignoring the updated value for the file-modification time.

22

1       7.       In a file server system having a first processor and a second processor for  
2       servicing client requests for access to a file, the first processor having a clock producing a  
3       clock time, and the second processor having a timer for measuring a time interval, a  
4       method comprising:

5               the second processor responding to a first request from a client for an  
6       asynchronous write to the file by obtaining the clock time from the clock of the first  
7       processor, beginning measurement of the time interval with the timer, performing a first  
8       asynchronous write operation with respect to the file, and using the clock time obtained  
9       from the clock of the first processor as a first file-modification time, the first file-  
10      modification time indicating a time of modification of the file by the first asynchronous  
11      write operation; and thereafter

12              the secondary processor responding to a second request from the client for an  
13      asynchronous write to the file by performing a second asynchronous write operation with  
14      respect to the file, and determining a second file-modification time that is a function of  
15      the clock time obtained from the clock of the first processor and the time interval  
16      measured by the timer, the second file-modification time indicating a time of  
17      modification of the file by the second asynchronous write operation.

18  
19      8.       The method as claimed in claim 7, wherein the file-modification time is a sum of  
20      the clock time having been obtained from the clock and the time interval measured by the  
21      timer.

22

1       9.       The method as claimed in claim 7, which includes:  
2               the second processor acknowledging the first request from the client for an  
3       asynchronous write to the file by returning to the client the first file-modified time for the  
4       file; and  
5               the second processor acknowledging the second request from the client for an  
6       asynchronous write to the file by returning to the client the second file-modified time for  
7       the file.

8  
9       10.      The method as claimed in claim 7, which includes the second processor  
10      responding to a request from the client to commit results of the second asynchronous  
11      write operation by sending the second file-modification time to the first processor.

12  
13     11.      In a file server system having a first processor and a second processor for  
14     servicing client requests for access to a file, the first processor having a clock producing a  
15     clock time, and the second processor having a timer for measuring a time interval, a  
16     method comprising:

17             the second processor responding to a first request from a client for an  
18     asynchronous write to the file by obtaining the clock time from the clock of the first  
19     processor, beginning measurement of the time interval with the timer, performing a first  
20     asynchronous write operation with respect to the file, and using the clock time obtained  
21     from the clock of the first processor as a first file-modification time, the first file-  
22     modification time indicating a time of modification of the file by the first asynchronous  
23     write operation; and thereafter

1           the second processor receiving from the first processor an updated value for the  
2   file-modification time, the second processor comparing the updated value for the file-  
3   modification time to the first file-modification time, and upon finding that the updated  
4   value is greater than the first file-modification time, the second processor resetting the  
5   timer; and thereafter

6           the second processor responding to a second request from the client for an  
7   asynchronous write to the file by performing a second asynchronous write operation with  
8   respect to the file, and determining a second file-modification time that is a sum of the  
9   updated value for the file-modification time and the time interval measured by the timer,  
10   the second file-modification time indicating a time of modification of the file by the  
11   second asynchronous write operation.

12

13   12.   In a file server system having a primary processor managing metadata of a file,  
14   and a secondary processor responding to requests from a client for access to the file, the  
15   primary processor having a clock producing a clock time, and the secondary processor  
16   having a timer for measuring a time interval, a method comprising:

17           the secondary processor responding to a first asynchronous write request from the  
18   client for writing to the file by obtaining attributes of the file and the clock time from the  
19   primary processor, storing the attributes of the file in a cache local to the secondary  
20   processor and using the file attributes to perform a first asynchronous write operation  
21   with respect to the file, and beginning measurement of the time interval with the timer,  
22   and thereafter

1           the secondary processor responding to a second asynchronous write request from  
2           the client for writing to the file by using the attributes of the file in the cache local to the  
3           secondary processor to perform a second asynchronous write operation with respect to  
4           the file, and determining a file-modification time that is a function of the clock time  
5           having been obtained from the clock of the primary processor and the interval measured  
6           by the timer, the file-modification time indicating a time of modification of the file by the  
7           second asynchronous write operation.

8  
9       13.     The method as claimed in claim 12, wherein the file-modification time is a sum of  
10       the clock time having been obtained from the clock and the time interval measured by the  
11       timer.

12  
13      14.     The method as claimed in claim 12, which includes:

14           the secondary processor acknowledging the second asynchronous write request  
15           from the client by returning to the client the file-modification time as the time when the  
16           file was modified by the second asynchronous write operation.

17  
18      15.     The method as claimed in claim 12, which includes:

19           the secondary processor responding to a request from the client to commit results  
20           of the second asynchronous write operation by sending a flush request to the primary  
21           processor, the flush request including the file-modification time.

1     16.     The method as claimed in claim 15, which includes the primary processor sending  
2     the file-modification time to another client caching attributes for the file.

3  
4     17.     The method as claimed in claim 12, which includes the secondary processor  
5     receiving from the primary processor an updated value for the file-modification time after  
6     the secondary processor has completed the second asynchronous write operation, the  
7     secondary processor comparing the updated value for the file-modification time to the  
8     last value for the file-modification time determined by the secondary processor, and upon  
9     finding that the updated value for the file-modification time is greater than the last value  
10    for the file-modification time determined by the secondary processor, the secondary  
11    processor resetting the timer, and using the updated value for the file-modification time in  
12    lieu of the clock time having been obtained from the primary processor, and using the  
13    updated value for the file-modification time as the most recent value of the file-  
14    modification time.

15  
16    18.     The method as claimed in claim 12, which includes the secondary processor  
17    receiving from the primary processor an updated value for the file-modification time after  
18    the secondary processor has completed the second asynchronous write operation, the  
19    secondary processor comparing the updated value for the file-modification time to the  
20    last value for the file-modification time determined by the secondary processor, and upon  
21    finding that the updated value for the file-modification time is less than the last value for  
22    the file-modification time determined by the secondary processor, the secondary  
23    processor ignoring the updated value for the file-modification time.

1

2     19.     In a network file server having a plurality of data mover computers for servicing  
3     client requests for access to a file, and a cached disk array for storing data of the file, the  
4     data mover computers being coupled to the cache disk array for accessing the data of the  
5     file, the data mover computers including a primary data mover computer managing  
6     metadata of the file, and a secondary data mover computer that requests metadata of the  
7     file from the primary data mover computer, the primary data mover computer having a  
8     clock producing a clock time, and the secondary data mover computer having a timer for  
9     measuring a time interval, a method comprising:

10           the secondary data mover computer responding to a first asynchronous write  
11     request from a client for writing to the file by obtaining attributes of the file and the clock  
12     time from the primary data mover computer, storing the attributes of the file in a cache  
13     local to the secondary data mover computer and using the file attributes to perform a first  
14     asynchronous write operation with respect to the file, and using the clock time as a first  
15     file-modification time indicating a time of modification of the file by the first  
16     asynchronous write operation; and thereafter

17           the secondary data mover computer responding to a second asynchronous write  
18     request from the client for writing to the file by using the attributes of the file in the cache  
19     local to the secondary data mover computer to perform a second asynchronous write  
20     operation with respect to the file, and determining a second file-modification time that is  
21     a function of the clock time having been obtained from the primary data mover and the  
22     time interval measured by the timer, the second file-modification time indicating a time  
23     of modification of the file by the second asynchronous write operation.



1

2     20.     The method as claimed in claim 19, wherein the second file-modification time is a  
3     sum of the clock time having been obtained from the primary data mover and the time  
4     interval measured by the timer.

5

6     21.     The method as claimed in claim 19, wherein:

7             the secondary data mover computer uses the clock time as a first file-modification  
8     time by acknowledging the first asynchronous write request from the client by returning  
9     to the client the clock time as the time when the file was modified by the first  
10    asynchronous write operation, and

11            the secondary data mover computer acknowledges the second asynchronous write  
12    request from the client by returning to the client the second file-modification time as the  
13    time when the file was modified by the second asynchronous write operation.

14

15    22.     The method as claimed in claim 19, which includes:

16             the secondary data mover computer responding to a request from the client to  
17    commit results of the second asynchronous write operation by sending a flush request to  
18    the primary data mover computer, the flush request including the second file-  
19    modification time.

20

21    23.     The method as claimed in claim 22, which includes the primary data mover  
22    computer sending the second file-modification time to another client caching attributes  
23    for the file.

1

2     24.     The method as claimed in claim 19, which includes the secondary data mover  
3     computer receiving from the primary data mover computer an updated value for the file-  
4     modification time for the file after the secondary data mover computer has completed the  
5     first asynchronous write operation, the secondary data mover computer comparing the  
6     updated value for the file-modification time for the file to the last value determined by  
7     the secondary data mover for the file-modified time for the file, and upon finding that the  
8     updated value for the file-modification time for the file is greater than the last value  
9     determined by the secondary data mover for the file-modified time for the file, the  
10    secondary data mover computer resetting the timer, using the updated value for the file-  
11    modification time in lieu of the clock time having been obtained from the primary data  
12    mover computer, and using the updated value for the file-modification time for the file as  
13    the most recent value for the file-modification time for the file.

14

15    25.     The method as claimed in claim 19, which includes the secondary data mover  
16    computer receiving from the primary data mover computer an updated value for the file-  
17    modification time for the file after the secondary data mover computer has completed the  
18    first asynchronous write operation, the secondary data mover computer comparing the  
19    updated value for the file-modification time for the file to the last value determined by  
20    the secondary data mover for the file-modified time for the file, and upon finding that the  
21    updated value for the file-modification time for the file is less than the last value  
22    determined by the secondary data mover for the file-modified time for the file; the  
23    secondary data mover computer ignoring the updated value for the file-modification time.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23

26. A file server system having a clock for producing a clock time and a processor for servicing client requests for access to a file, the processor having a timer for measuring a time interval;
- the processor being programmed for obtaining the clock time from the clock, and beginning measurement of the time interval with the timer, and
- the processor being programmed for responding to a request from a client for an asynchronous write to the file by performing an asynchronous write operation with respect to the file, and determining a file-modification time that is a function of the clock time having been obtained from the clock and the time interval measured by the timer, the file-modification time indicating a time of modification of the file by the asynchronous write operation.
27. The file server system as claimed in claim 26, wherein the file-modification time is a sum of the clock time having been obtained from the clock and the time interval measured by the timer.
28. The file server system as claimed in claim 26, wherein the processor is programmed to acknowledging the request from the client for an asynchronous write to the file by returning to the client the file-modification time.
29. The file server system as claimed in claim 26, wherein the processor is programmed for receiving an updated value for the file-modification time after the

1 processor has determined a value for the file-modification time, comparing the updated  
2 value for the file-modification time to the value that the processor has determined for the  
3 file-modification time, and upon finding the updated value for the file-modification time  
4 is greater than the value that the processor has determined for the file-modification time,  
5 resetting the timer and using the updated value for the file-modification time in lieu of the  
6 clock time having been obtained from the clock.

7  
8 30. The file server system as claimed in claim 29, wherein the processor is  
9 programmed for storing the clock time having been obtained from the clock in a memory  
10 location local to the processor, and for using the updated value for the file-modification  
11 time in lieu of the clock time having been obtained from the clock by replacing the clock  
12 time stored in the memory local to the processor with the updated value for the file-  
13 modification time.

14  
15 31. The file server system as claimed in claim 26, wherein the processor is  
16 programmed for receiving an updated value for the file-modification time after the  
17 secondary processor has determined a value for the file-modification time, comparing the  
18 updated value to the value that the processor has determined for the file-modification  
19 time, and ignoring the updated value for the file-modification time upon finding the  
20 updated value for the file-modification time is less than the value that the secondary  
21 processor has determined for the file-modification time.

1     32.     A file server system comprising:  
2             a first processor and a second processor for servicing client requests for access to  
3     a file, the first processor having a clock for producing a clock time, and the second  
4     processor having a timer for measuring a time interval;  
5             the second processor being programmed for responding to a first request from a  
6     client for an asynchronous write to the file by obtaining the clock time from the clock of  
7     the first processor, beginning measurement of the time interval with the timer,  
8     performing a first asynchronous write operation with respect to the file, and using the  
9     clock time obtained from the clock of the first processor as a first file-modification time,  
10    the first file-modification time indicating a time of modification of the file by the first  
11    asynchronous write operation; and the second processor being programmed for  
12    responding to a second request from the client for an asynchronous write to the file by  
13    performing a second asynchronous write operation with respect to the file, and  
14    determining a second file-modification time that is a function of the clock time obtained  
15    from the clock of the first processor and the time interval measured by the timer, the  
16    second file-modification time indicating a time of modification of the file by the second  
17    asynchronous write operation.

18  
19    33.     The file server system as claimed in claim 32, wherein the second file-  
20    modification time is a sum of the clock time obtained from the clock and the time interval  
21    measured by the timer.

22

1     34.     The file server system as claimed in claim 32, wherein:  
2             the second processor is programmed to use the clock time obtained from the clock  
3     of the first processor as a first file-modification time by acknowledging the first request  
4     from the client for an asynchronous write to the file by returning to the client the clock  
5     time obtained from the clock of the first processor as the time when the file was modified  
6     by the first asynchronous write operation, and  
7             the second processor is programmed to acknowledge the second request from the  
8     client for an asynchronous write to the file by returning to the client the second file-  
9     modification time as the time when the file was modified by the second asynchronous  
10    write operation.

11

12    35.     The file server system as claimed in claim 32, wherein the second processor is  
13    programmed for responding to a request from the client to commit results of the second  
14    asynchronous write operation by sending the second file-modification time to the first  
15    processor.

16

17    36.     A file server system comprising:  
18             a first processor and a second processor for servicing client requests for access to  
19    a file, the first processor having a clock for producing a clock time, and the second  
20    processor having a timer for measuring a time interval;  
21             the second processor being programmed for responding to a first request from a  
22    client for an asynchronous write to the file by obtaining the clock time from the clock of  
23    the first processor, beginning measurement of the time interval with the timer, and

1 performing a first asynchronous write operation with respect to the file, and using the  
2 clock time obtained from the clock of the first processor as a first file-modification time  
3 indicating a time of modification of the file by the first asynchronous write operation; and  
4 the second processor being programmed for receiving from the first processor an  
5 updated value for the file-modification time, for comparing the updated value to the first  
6 file-modification time, and upon finding that the updated value is greater than the first  
7 file-modification time, for resetting the timer; and  
8 the second processor being programmed to respond to a second request from the  
9 client for an asynchronous write to the file by performing a second asynchronous write  
10 operation with respect to the file, and determining a second file-modification time that is  
11 a sum of the updated value for the file-modification time and the time measured by the  
12 timer, the second file-modification time indicating a time of modification of the file by  
13 the second asynchronous write operation.

14

15 37. A file server system comprising:

16 a primary processor managing metadata of a file, and a secondary processor  
17 responding to requests from a client for access to the file, the primary processor having a  
18 clock for producing a clock time, and the secondary processor having a timer for  
19 measuring a time interval;

20 the secondary processor being programmed for responding to a first asynchronous  
21 write request from the client for writing to the file by obtaining attributes of the file and  
22 the clock time from the primary processor, storing the attributes of the file in a cache  
23 local to the secondary processor and using the file attributes to perform a first

1 asynchronous write operation with respect to the file, and beginning measurement of the  
2 time interval with the timer; and

3 the secondary processor being programmed for responding to a second  
4 asynchronous write request from the client for writing to the file by using the attributes of  
5 the file in the cache local to the secondary processor to perform a second asynchronous  
6 write operation with respect to the file, and determining a file-modification time that is a  
7 function of the clock time having been obtained from the clock of the primary processor  
8 and the time interval measured by the timer, the file-modification time indicating a time  
9 of modification of the file by the second asynchronous write operation.

10

11 38. The file server system as claimed in claim 37, wherein the file-modification time  
12 is a sum of the clock time having been obtained from the primary processor and the time  
13 interval measured by the timer.

14

15 39. The file server system as claimed in claim 37, wherein:

16 the secondary processor is programmed for acknowledging the second  
17 asynchronous write request from the client by returning to the client the file-modification  
18 time as the time when the file was modified by the second asynchronous write operation.

19

20 40. The file server system as claimed in claim 37, wherein the secondary processor is  
21 programmed for responding to a request from the client to commit results of the second  
22 asynchronous write operation by sending a flush request to the primary processor, the  
23 flush request including the file-modification time.



1

2     41.     The file server system as claimed in claim 40, wherein the primary processor is  
3     programmed to send the file-modification time to other clients caching attributes for the  
4     file.

5

6     42.     The file server system as claimed in claim 37, wherein the secondary processor is  
7     programmed for receiving from the primary processor an updated value for the file-  
8     modification time after the secondary processor has completed the second asynchronous  
9     write operation, for comparing the updated value for the file-modification time to the last  
10    value for the file-modification time determined by the secondary processor, and upon  
11    finding that the updated value for the file-modification time is greater than the last value  
12    for the file-modification time determined by the secondary processor, for resetting the  
13    timer, and using the updated value for the file-modification time in lieu of the clock time  
14    having been obtained from the primary processor, and using the updated value for the  
15    file-modification time as the most recent value of the file-modification time.

16

17    43.     The file server system as claimed in claim 37, wherein the secondary processor is  
18    programmed for receiving from the primary processor an updated value for the file-  
19    modification time after the secondary processor has completed the second asynchronous  
20    write operation, for comparing the updated value for the file-modification time to the last  
21    value for the file-modification time determined by the secondary processor, and upon  
22    finding that the updated value for the file-modification time is less than the last value for

1 the file-modification time determined by the secondary processor, for ignoring the  
2 updated value for the file-modification time.

3  
4 44. A network file server comprising:

5 a plurality of data mover computers for servicing client requests for access to a  
6 file, and a cached disk array for storing data of the file, the data mover computers being  
7 coupled to the cached disk array for accessing the data of the file, the data mover  
8 computers including a primary data mover computer programmed for managing metadata  
9 of the file, and a secondary data mover computer programmed for requesting metadata of  
10 the file from the primary data mover computer, the primary data mover computer having  
11 a clock for producing a clock time, and the secondary data mover computer having a  
12 timer for measuring a time interval;

13 the secondary data mover computer being programmed for responding to a first  
14 asynchronous write request from a client for writing to the file by obtaining attributes of  
15 the file and the clock time from the primary data mover computer, storing the attributes  
16 of the file in a cache local to the secondary data mover computer and using the file  
17 attributes to perform a first asynchronous write operation with respect to the file,  
18 beginning measurement of the time interval with the timer, and using the clock time as a  
19 first file-modification time, the first file-modification time indicating a time of  
20 modification of the file by the first asynchronous write operation; and

21 the secondary data mover computer being programmed for responding to a second  
22 asynchronous write request from the client for writing to the file by using the attributes of  
23 the file in the cache local to the secondary data mover computer to perform a second

1 asynchronous write operation with respect to the file, and determining a second file-  
2 modification time that is a function of the clock time having been obtained from the  
3 primary data mover and the time interval measured by the timer, the second file-  
4 modification time indicating a time of modification of the file by the second  
5 asynchronous write operation.

6

7 45. The network file server as claimed in claim 44, wherein the second file-  
8 modification time is a sum of the clock time having been obtained from the primary data  
9 mover and the time interval measured by the timer.

10

11 46. The network file server as claimed in claim 44, wherein:

12 the secondary data mover computer is programmed for using the clock time as a  
13 first file-modification time by acknowledging the first asynchronous write request from  
14 the client by returning to the client the clock time as the time when the file was modified  
15 by the first asynchronous write operation, and

16 the secondary data mover computer is programmed for acknowledging the second  
17 asynchronous write request from the client by returning to the client the second file-  
18 modification time as the time when the file was modified by the second asynchronous  
19 write operation.

20

21 47. The network file server as claimed in claim 44, wherein the secondary data mover  
22 computer is programmed for responding to a request from the client to commit results of

1 the second asynchronous write operation by sending a flush request to the primary data  
2 mover computer, the flush request including the second file-modification time.

3

4 48. The network file server as claimed in claim 47, wherein the primary data mover  
5 computer is programmed for sending the second file-modification time to other clients  
6 caching attributes for the file.

7

8 49. The network file server as claimed in claim 44, wherein the secondary data mover  
9 computer is programmed for receiving from the primary data mover computer an updated  
10 value for the file-modification time for the file after the secondary data mover computer  
11 has completed the first asynchronous write operation, for comparing the updated value  
12 for the file-modification time for the file to last value determined by the secondary data  
13 mover for the file-modified time for the file, and upon finding that the updated value for  
14 the file-modification time for the file is greater than the last value determined by the  
15 secondary data mover for the file-modified time for the file, for resetting the timer, using  
16 the updated value for the file-modification time for the file in lieu of the clock time  
17 having been obtained from the primary data mover computer, and using the updated  
18 value for the file-modified time for the file as the most recent value for the file-  
19 modification time for the file.

20

21 50. The network file server as claimed in claim 44, wherein the secondary data mover  
22 computer is programmed for receiving from the primary data mover computer an updated  
23 value for the file-modification time for the file after the secondary data mover computer

1 has completed the first asynchronous write operation, for comparing the updated value  
2 for the file-modification time for the file to last value determined by the secondary data  
3 mover for the file-modified time for the file, and upon finding that the updated value for  
4 the file-modification time for the file is less than the last value determined by the  
5 secondary data mover for the file-modified time for the file, for ignoring the updated  
6 value for the file-modification time for the file.